EXHIBIT B

Engineering Specification	Security Notation **Confidential**	Spec No. **03-220-0014-03** Tentative P/N	Rev Ltr 03
Type **Analysis**	Class **C**	Initial Release Date 8/09/03	
Division DPA - 5182 / Dept No. 1075232	Model No.	Part No.	

Title

passage [redacted] Trial Phase-3 VOD Architecture Study

Prepared By L.M. Pedlow, D. Agnihotri — Date 10/16/03	Approved by Engr Mgr — Date	Approved By — Date
Approved For — Date	Approved For — Date	Approved For — Date

EXHIBIT B

FOR REVISIONS, SEE CR-2. REVISION RECORD FOLLOWS DOCUMENT CONTENTS

# Engineering Specification Revision History	Security Notation **Confidential**	Spec No. **03-220-0014-03**	Rev Ltr 03

See First Page for Proprietary or Data Rights Notations

Rev Ltr	Description	Author	Date and Approval
00	Preliminary Draft	L.M. Pedlow	08-09-03
01	1st Revision	L.M. Pedlow	08-27-03
02	2nd Revision – Added information on VOD server PSI	L.M. Pedlow	09-02-03
03	3rd Revision Inserted Figure 3 and clarified pre-processing	L.M. Pedlow	10-16-03

Contents

Figures

passage VOD Architecture Study

1 Scope and Overview

This document provides an overview of the various video on demand (**VOD**) architectures that warrant consideration as part of integrating open standard, alternative conditional access architecture on an existing, incumbent digital cable television network based upon either Motorola or Scientific Atlanta systems. The intent of this document is to highlight the issues associated with each architecture and to compare the positive and negative aspects of each architecture so that the MSO can be better prepared to make informed decisions regarding the best approach for their unique business and network needs.

Depending upon the specific VOD architecture in question, Passage™ processing technology may be actively or passively employed. In most cases, Passage™ processing technology, as used in the broadcast portions of digital cable systems, is not employed at all in the distribution of VOD content. For continuity, it is presumed that the reader is already familiar with the Passage™ processing concept as applied to broadcast transport streams and therefore a discussion of the technology will not be presented in this document. Literature providing a complete background in Passage™ broadcast stream processing may be found in the cited references.

1.1 Overview

The purpose of the Passage™ initiative, promoted by Sony, is to facilitate competition in an open marketplace by providing a means for MSOs to deploy non-legacy headend equipment, subscriber devices and services on their existing legacy networks. In the USA, these networks are supplied by either Motorola (former General Instrument) or Scientific Atlanta. These two companies at present constitute better than a 99% share of the US cable system market as turnkey system providers. The systems, by design, employ proprietary technology and interfaces precluding the introduction of non-incumbent equipment into the network. An MSO, once choosing one of these suppliers during conversion from an analog cable system to a digital cable system, faces a virtual monopoly when seeking suppliers for additional equipment as their subscriber base or service offering grows.

Before the Passage™ initiative, the only exit from this situation was to forfeit the considerable capital investment already made with the incumbent provider, due to the intentional incompatibility of equipment between the incumbent and other sources. One primary barrier to interoperability is in the area of conditional access systems, the heart of addressable subscriber management and revenue collection resources in a modern digital cable network.

Passage™ was developed to allow the independent coexistence of two or more conditional access systems on a single, common plant. Unlike other attempts to address the issue, the two systems operate with a common transport stream without any direct or indirect interaction between the conditional access systems. This process is discussed in detail in the cited references.

While Passage™ was originally targeted at broadcast transport streams in digital cable plants, the very nature of needing real-time conditional access and a "one to many" information model reduces the number of conceptual architectures to a small number and in general a single topology generally suits most applications.

As will be shown, VOD is a much more complex issue.

1.2 Reference Documents

Unless a specific version or revision number is included, the following references are non-specific, and the latest version applies. For Sony *Passage*™ specifications, these non-specific references are indicated by the "xx" instead of a revision number, for example, "02-212-0002-xx".

1. Sony – "*Passage*™ Processor Specification" (02-212-0002-xx)
2. Sony – "*Passage*™ Set-top Box Specification" (02-212-0006-xx)
3. Sony – "*Passage*™ Decoder Module Specification" (02-212-0007-xx)
4. Sony – "*Passage*™ EIS-to-PE Interface Control Document" (02-214-0001-xx)
5. Sony – "Passage™ Encoder Configuration Server to Encoder Communication ICD" (02-214-0002-xx)
6. Sony – "*Passage*™ Auxiliary PAT Identifier" (02-214-0004-xx)
7. SCTE – "*OpenCAS*™ DVS278R1" (Version July 31, 2000)
8. *Head-end Implementation of DVB Simulcrypt*, Draft ETSI TS103 197 V1.3.1 (02-06) TM2117R3
9. *Information Technology — Generic coding of moving pictures and associated audio information: systems — Part 1: Systems*, ISO/IEC 13818-1, ISO, 4/02.
10. *Digital Multi-Programme Systems for Television, Sound and Data Services for Cable Distribution*, ITU-T J.83, International Telecommunications Union, 04/97
11. *Cable Networks for Television Signals, Sound Signals and Interactive Services Part 9: Interfaces for CATV/SMATV Headends and Similar Professional Equipment for DVB/MPEG2 Transport Streams*, EN 50083-9:1998, CENELEC, 1998

1.3 Acronyms and Abbreviations

ASI	Asynchronous Serial Interface
CA	Conditional Access
CASID	Conditional Access System Identifier
CPE	Customer Premises Equipment
DHEI	Digital Headend Extended Interface
ECM	Entitlement Control Message
EPG	Electronic Program Guide
GOP	Group of Pictures (MPEG)
MPEG	Moving Pictures Experts Group
MSO	Multiple System Operator

PAT	Program Allocation Table
PID	Packet Identifier
PMT	Program Map Table
PSI	Program Specific Information
QAM	Quadrature Amplitude Modulation
RAM	Random Access Memory
SAN	Storage Area Network
VOD	Video on Demand

2 Comparative VOD Architectures

The decision on a particular VOD architecture is the result of the interaction between a complex set of both independent and dependent variables, providing a solution to an equation of state. Some of the variables are fixed directly as a result of choices by the MSO. Others are constrained by factors such as the existing incumbent system, location, size, available capital and ROI requirements.

A generalized VOD system, shown in Figure 1, contains some or all of the following elements:

- Content aggregation
- Asset management
- Content distribution (SAN)
- Video server module(s)
- Session management
- Transaction management
- Billing system interface
- EPG server interface or VOD catalog server
- Transport router/switch fabric
- Stream encryption device(s)
- QAM modulators/upconverters

One aspect of VOD that has become a "signature" feature is the support of "trick modes". These are operational modes invoked by the session client that mimic a traditional VCR or DVD player and includes fast forward, rewind, pause, suspend (stop), slow motion, etc. Trick modes are typically implemented through the creation of multiple files containing a subset of the original content. Typically, these subfiles consist of only I-frames, since they are stand-alone whole pictures[1]

[1] ISO/IEC 13818-2, section 6.1.1.7

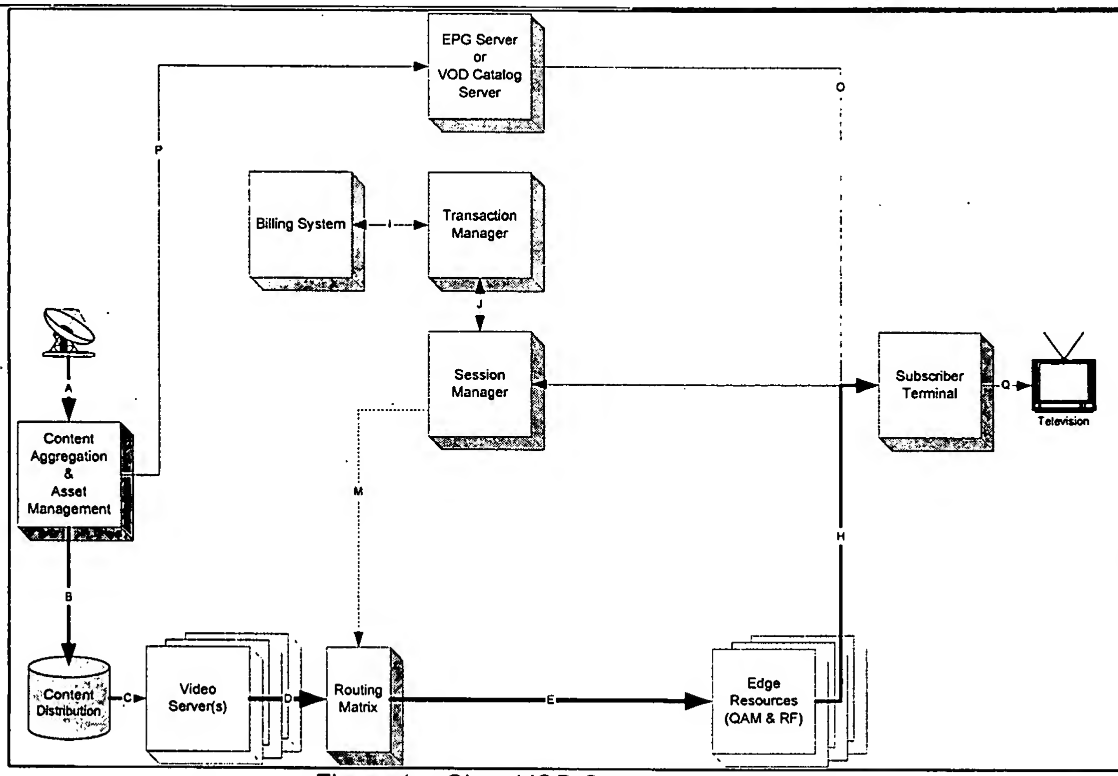

Figure 1 – *Clear VOD System Architecture*

A file containing only I-frames extracted from the original content affords the ability to have accelerated playback, since typical GOP structures have only one frame in 10 or 20 as an I-frame. If the I-frame files are played at normal rates (1 frame per 33 mS) the pictures will appear to the viewer to sequence at a 10x or 20x rate, though the actual data rate is the same as the original content. If the I-frame sequence is reversed in the file, the motion will appear to run backwards. This is the method used to implement fast forward & rewind.

By attaching an index count to match the I-frames in the original content file to the duplicated I-frames stored in the associated subfiles, a method is provided to allow immediate transition from normal speed forward play to fast forward or rewind (Figure 2). In operation the video server plays the selected content file and upon subscriber selection of a trick mode (or vice versa) the server notes the index value of the closest I-frame and then opens the appropriate associated subfile and moves to the I-frame in subfile with the same corresponding index. The video server treats all stream content (main file or subfiles) the same and always spools the MPEG packets to the outgoing transport stream at the same constant bit rate. It is through this clever method that trick modes are typically implemented on a slotted, session based system without the encumbrance of additional, dynamic bit rate issues.

A key function of the VOD server, in addition to origination of session A/V content, is the creation of the associated, session specific PSI. This information is a departure from the broadcast model in that the PSI is extremely dynamic. The content of the PAT and subordinate PMTs change whenever a new session is started or ended. In the broadcast world, the PSI changes very seldom because the PSI tables reflect only the structure of the transport multiplex, not the actual A/V content carried within.

The VOD video server must dynamically assign a new session to an existing, available "slot" in an outgoing transport multiplex. The slot is denoted by the MPEG program number and in many cases, the combination of which transport stream (TSID) and program number determine at the service level a unique session and the routing that occurring as a result. Edge devices typically are not configured dynamically. The routing of content appearing on a particular input port to a specific QAM carrier at the output is determined through a preconfigured, static assignment of TSID/input port and program number mapping to specific QAM resources in the device. This same mapping information is also loaded in the VOD system so that once a session is requested by and authorized for a specific subscriber, a solution to a routing matrix can be determined to find the appropriate VOD server and QAM transport serving the requestor. This solution also considers dynamic issues such as which servers the requested asset is loaded upon, and server loading/available slots in addition to the simpler, static solution to finding the first possible path to the requestor.

In addition to solving the routing matrix and provisioning the session with PIDs and PSI appropriate to follow the intended route, elements of the same information (program ID and QAM frequency) must also be communicated to the session client at the subscriber's premises so that the requested stream can be received and presented to the subscriber.

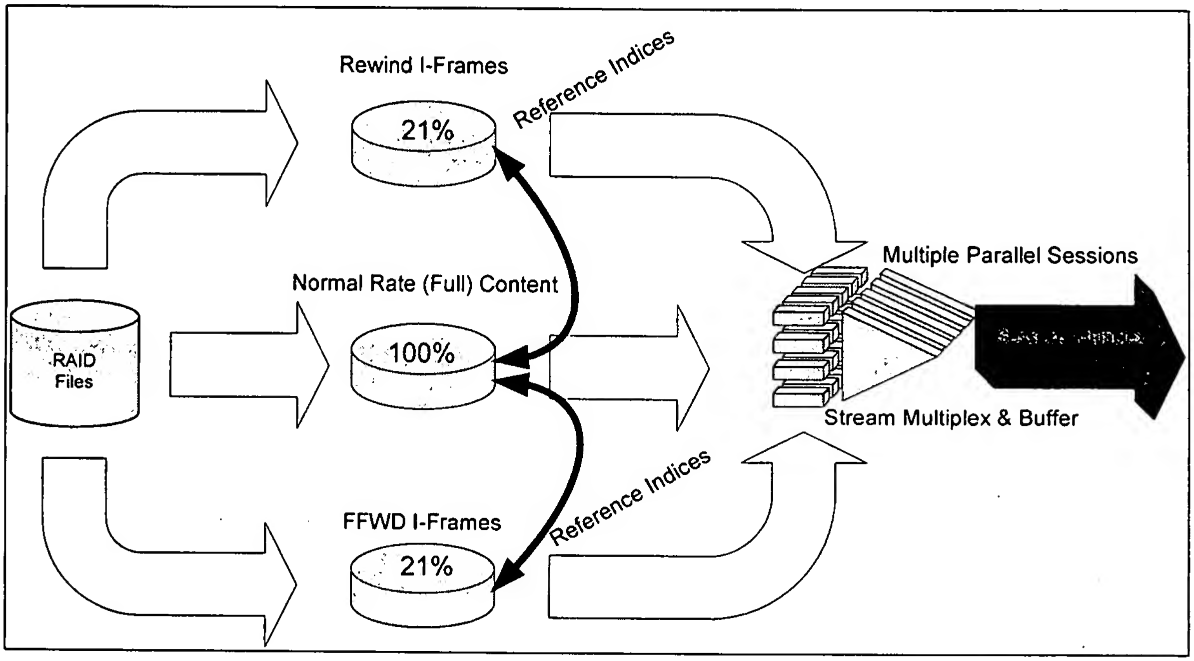

Figure 2 – *Clear VOD Session Content Flow*

2.1 Clear VOD Distribution

The simplest VOD implementation is a clear VOD distribution system, i.e. one that contains no encryption. While not providing any safekeeping of what might be considered the entertainment medium's most valuable properties, namely current feature films, etc., clear VOD avoids many of the issues that the incumbent cable system providers to date have not adequately addressed and that introduction of a second, alternative CA system complicates even further still.

A real world example:

- Compressed video data rate: 3Mbit/S
- Movie length: 120 minutes (2 Hrs)
- I-frame overhead: 17%
- Total storage required for the video portion of a <u>single</u>, clear (unencrypted) copy of a film: **3.618GBytes**

2.2 Pre-Encrypted VOD Distribution

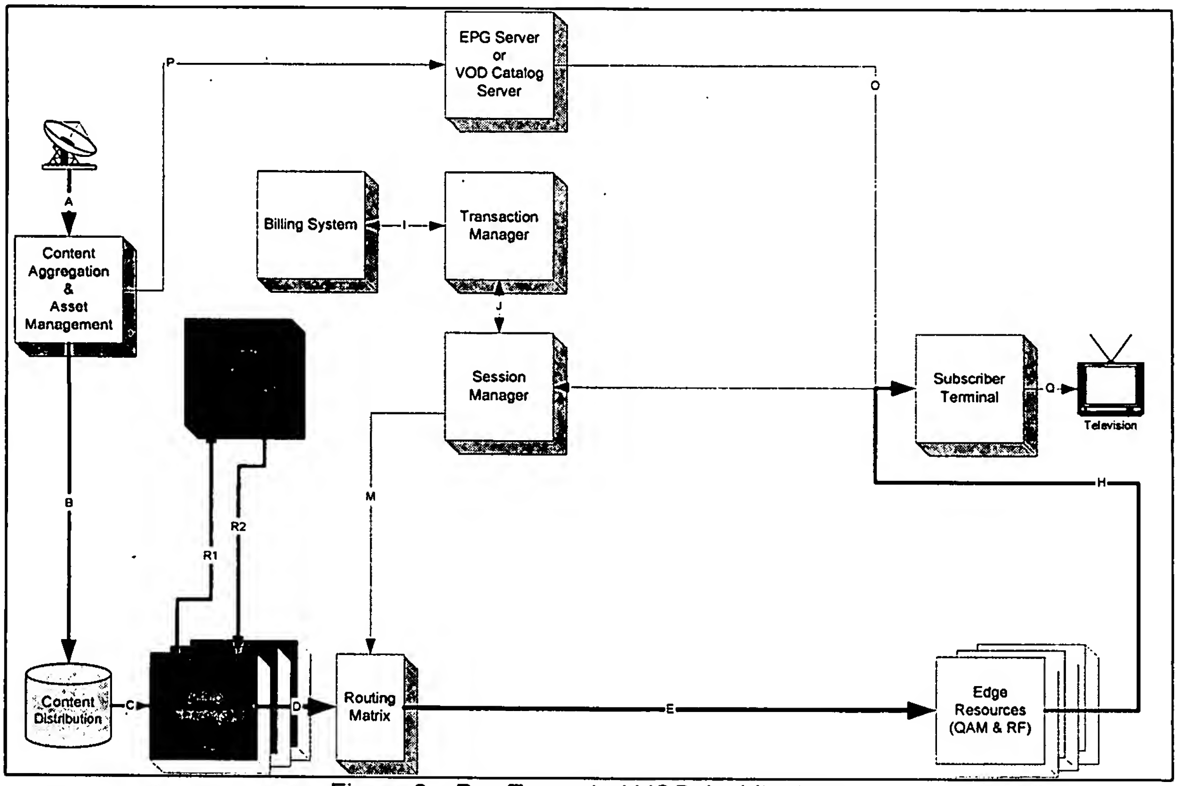

Figure 3 - *Pre-Encrypted VOD Architecture*

Pre-encrypted VOD systems are architecturally identical to clear VOD distribution systems. The difference between the two is that on pre-encrypted systems there is pre-processing of the content prior to storage in the VOD system to provide safekeeping of content during the storage and distribution phases. Data security is implemented through storage of previously encrypted content within the video server(s). While the clear VOD system contains directly viewable MPEG or other compressed A/V content on the server, the pre-encrypted model stores this same content in a form that is only decipherable using a properly entitled subscriber terminal.

The pre-encryption process can be performed by the MSO at the time of deployment on the VOD system, prior to loading into the storage area network (SAN) used to propagate content to all of the video servers in the MSO's system. Alternatively, the encryption may be performed prior to receipt of the content by the MSO at an external service bureau, content aggregator or by the distributor or studio. In this case, the content is theoretically secured throughout the distribution phase, storage phase and transmission to subscriber for display on an authorized device. The use of pre-encryption prior to distribution of content to the MSO potentially adds the complexity of entitlement distribution, separate from the content distribution, for installation on the VOD transaction manager to allow bone fide subscribers to decrypt the purchased content.

Most pre-encrypted VOD architectures share one or more of the following common drawbacks:

- Additional handling of new content is necessary to perform the pre-encryption prior to loading into the server, either by the MSO or service bureau.
- Coordination and/or distribution required for entitlements matching the access criteria used to encrypt the content stored in the server.
- Limited "shelf life" of the encryption keys used to secure the stored content, rendering decryption impossible at a later date.
- Incapability of present VOD video servers to load pre-encrypted streams
- Incompatibility of pre-encrypted streams with present methods supporting trick mode play (fast-forward & rewind) on screen.
- One common key is used for all sessions accessing a particular program and it remains the same for the duration of time the content is in inventory on the server.
- According to MSOs familiar with the subject, pre-encrypted VOD streams are unsupported by present Scientific-Atlanta conditional access technologies.

The issue regarding trick play and pre-encryption is based upon the concept that VOD servers currently expect clear content and then subsequently identify the I-frames and store or otherwise segregate them for access in fast-forward or fast rewind playback modes. If the stream is pre-encrypted prior to storage upon the server, it is impossible for the server to examine payloads to identify I-frames during the process of importation into the server to create trick mode files or indices. Many current systems will not accept streams for importation that are pre-encrypted.

2.2.1 Segregated Storage Pre-Encryption

Segregated storage is the simplest from a conceptual perspective, of all pre-encrypted VOD schemes. It is physically identical to the architecture of the clear VOD distribution system. The content is encrypted in its entirety (100%) and a separate copy of the complete feature is stored for each different conditional access format supported by the MSO. The organization and configuration of the system is such that when a subscriber initiates a session on the server, the stream files for the selected content containing the CA format appropriate to the specific equipment deployed at the subscriber's premises requesting the session are spooled and delivered. This method offers the lowest system complexity of any encrypted VOD system but suffers from the same issues common to other pre-encryption topologies, mentioned previously. In addition, a significant storage penalty (duplicate copies of the same movie) is incurred.

If one refers to the example movie scenario described in section 2.1, the same movie using 3.618GB of storage in the clear VOD state would require 7.236GBytes to store using segregated pre-encryption supporting two different CA systems.

Changes are required to the method employed by the VOD system for creating dynamic PSI data to implement this architecture supporting multiple CA systems. The VOD system session manager must be aware of which conditional access method is appropriate for a session requested by a specific subscriber. This information must in turn be transferred to the video server that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The video server must be cognizant of the conditional access resources (ECMs) for each program stored on the server and these must be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, must indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

2.2.2 Composite Storage Pre-Encryption

Composite storage is essentially the storage on the video server of a Passage™ processed stream that contains previously encrypted "critical packets" for two or more independent conditional access systems. The stream is prepared identically to the processing of a Passage™ broadcast stream, except that the resultant transport stream is recorded to a hard disk instead of being sent to a QAM modulator for HFC distribution to the requesting subscriber. As with other pre-encryption models, the content is encrypted by either the MSO at time of deployment on the VOD system, a third party service bureau or by the studios themselves, the latter two cases being prior to receipt of the content by the MSO.

The advantage of this architecture is the small additional overhead in content storage (typically 2% - 10%) traded for the support of multiple independent CA formats without replication of entire streams. The negative aspect, in addition to those mentioned previously and common to other pre-encryption topologies, is the vulnerability of the prepared Passage™ stream to corruption by downstream equipment containing transport remultiplexing functionality that is not specifically designed to maintain the integrity of the Passage™ process applied to the stream.

If one refers to the example movie scenario described in section 2.1, the same movie using 3.618GB of storage in the clear VOD state would require 3.690GBytes to store using composite storage pre-encryption supporting two different CA systems with a Passage™ density of 2%.

Changes are required to the method employed by the VOD system for creating dynamic PSI data to implement this architecture. The VOD system session manager must be aware of which conditional access method is appropriate for a session requested by a specific subscriber. This information must in turn be transferred to the video server that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The video server must be cognizant of the conditional access resources (ECMs) for each program stored on the server and these must be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, must indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

Likewise, the video server must dynamically allocate PIDs for the shadow packets associated with the respective audio and video component streams for each session. This information must also be included in the PSI sent in sessions requested by non-legacy clients. In total, eight different PIDs and corresponding data resources must be dynamically allocated and managed by the server for each session: PAT (one table common to all sessions, but modified for each), PMT, Primary Video, Primary Audio, Shadow Video, Shadow Audio, Legacy ECM and Alternative ECM. Six of these entities are stored in the embedded stream and require dynamic PID remapping for each session.

There is also the issue of what device to use in conjunction with performing the legacy encryption of the "critical" packets prior to storage on the VOD video server. If the legacy device is specially designed to process content destined for loading into a VOD video server, it is likely not to accept a Passage™ processed stream at its input. The content format specified for VOD servers requires a single program transport multiplex consisting of a single PAT entry, single PMT entry and service components, consisting of one audio and one video stream. The Passage™ shadow packets added in a composite Passage™ transport stream will likely prove problematic for a legacy VOD pre-encryption device. It is more probable that a device or process (since there are no real time requirements, an off-line process running on a PC or UNIX server may suffice) to process a candidate stream before passing through the legacy pre-encryptor and then post-encryption reconcile to extract only the encrypted "critical" packets for insertion into the VOD video server will be necessary. The algorithms and techniques for performing this manipulation are described in Sony Passage™ processing specifications and can be adapted to VOD applications for off-line work.

The VOD server will also require modification to allow introduction of streams consisting of multiple service elements (primary video, primary audio, shadow video, shadow audio) uniquely associated with a Passage™ transport. The present video servers only allow one each, primary video and audio, respectively. The quartet of data representing Passage™ processed A/V content must be managed as a indivisible set on the VOD video server.

Some additional bandwidth efficiencies may be obtained if, at the edge device, Passage™ shadow packets are removed from the composite streams in sessions serving legacy clients. Similarly, the edge device, if Passage™ aware, could reinsert the Passage™ shadow packets embedded in the stored stream in place of the legacy encrypted packets on the original program PID. These improvements would result in no carriage overhead for support of multiple conditional access systems on a single transport.

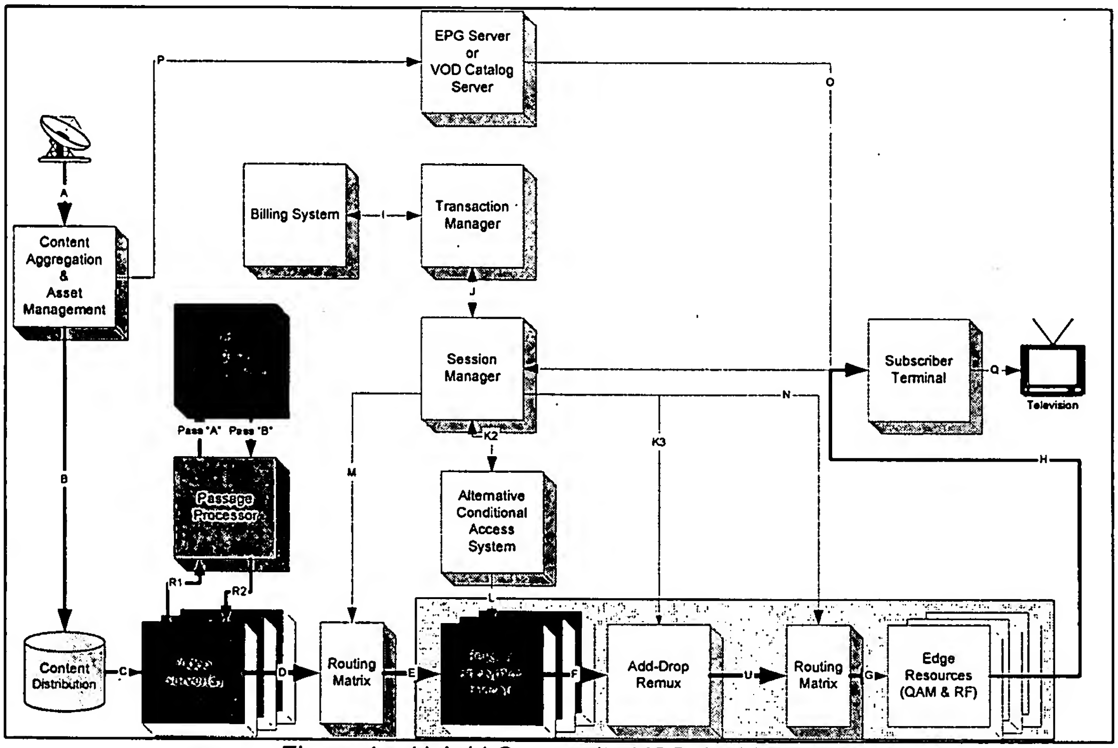

Figure 4 – *Hybrid Composite VOD Architecture*

2.2.3 Hybrid Composite Storage Pre-Encryption

Hybrid composite storage is a variant of the composite storage concept, but incorporates elements of session-based encryption for implementing the alternative conditional access encryption. In this scenario, the legacy "critical" packets, comprising 2-10% of the total content, are pre-encrypted by the incumbent conditional access system using Passage™ technology for managing the process. The duplicate copy of "critical" packets, which per the Passage™ process are located on previously unused PIDs, is left unencrypted. This latter aspect is the departure from the composite storage scenario described above. The composite stream consisting of unencrypted non-critical packets, legacy encrypted "critical" packets on the original service PIDs and an unencrypted, duplicate copy of the "critical" packets on alternate service PIDs is stored on the video server as a single stream.

Upon playback to a subscriber session, if the session is destined for a legacy STB, the existing paradigm for pre-encrypted content is followed and no special action is taken. The stream is routed through a device capable of performing encryption using the alternative conditional access system, but the session manager does not provision the device to perform encryption on elements of the stream and it is sent directly to the requesting subscriber. To maintain security of the outgoing stream and to reduce the bandwidth of the session for legacy sessions, the stream is processed through an add-drop remultiplexer and the clear "critical" content on alternate service PIDs are removed from the outgoing transport. It is likely that the device that performs encryption using the alternative conditional access system also contains the add-drop multiplexer capability.

If the session is destined for a non-legacy STB, the stream is routed through a device capable of performing encryption using the alternative conditional access system and only the "critical" packets on alternate service PIDs (previously in the clear) are encrypted using the alternative conditional access system, as provisioned by the session manager. Figure 4 shows the topology of a proposed VOD system employing hybrid composite storage pre-encryption.

Some additional bandwidth efficiencies may be obtained for these non-legacy sessions, if the edge device is Passage™ aware, by reinserting the Passage™ shadow packets embedded in the stored stream, now encrypted, in place of the legacy encrypted packets on the original program PID. This improvement would result in no carriage overhead for support of multiple conditional access systems on a single transport.

A preprocessor would be required to perform selective encryption of content to be loaded onto the video server. This could simply be a reapplication of Passage™ processors currently under development by Passage™ licensees. A modified file protocol would be required to allow the video server to import and associate these files. Either the preprocessor or the video server could perform the indexing. An alternate instantiation would be to perform all Passage™ pre-processing within the video server itself. This could be accomplished by modifying the video server application of to add a pre-processor task as a separate executable, called by the server during the process to prepare content for pre-encryption. Obviously, these relationships need further clarification and codification in a series of system specifications.

Changes are required to the method employed by the VOD system for creating dynamic PSI data to implement this architecture. The VOD system session manager must be aware of which conditional access method is appropriate for a session requested by a specific subscriber. This information must in turn be transferred to the video server that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The video server must be cognizant of the conditional access resources (ECMs) for each program stored on the server and these must be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, must indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

Likewise, the video server must dynamically allocate PIDs for the shadow packets associated with the respective audio and video component streams for each session. This information must also be included in the PSI sent in sessions requested by non-legacy clients. Just like in the more general composite storage architecture discussed in the previous section, the video server must manage multiple resources and PIDs. The hybrid topology reduces the unique entities by one from eight to seven: there is no alternative ECM PID or data resource in the stored composite stream. This information will be added later in a downstream device providing the alternative conditional access encryption for those sessions destined for decoding upon a non-legacy client.

2.2.4 Re-Encrypted Distribution

A hybrid approach is offered in a re-encrypted distribution architecture. This topology leverages the paradigms established for pre-encrypted content preparation, storage, management, etc. but adds support for session based encryption for the alternative conditional access systems added to an existing incumbent system. Referring to Figure 5, a legacy decryption device is added to the transport stream path exiting the VOD video server. After the decryption device, the transport stream passes through a contemporary session based encryption device. The VOD session manager, on a session-by-session basis, determines which sessions will pass through the decryption device intact and be modulated and transmitted to the subscriber unaltered. This path (T) between the routing matrices preserves the pre-encrypted content and delivers it to subscribers having legacy equipment. Alternatively, the VOD system session manager, through interaction with both CA systems, can both actuate the decryption device and activate session based encryption system for a particular session, supporting subscribers with non-legacy equipment at their premises.

The advantage of this architecture is the ability to support pre-encryption on legacy systems not presently supporting session-based encryption, while providing the ability to deliver session based encryption for the alternative CA system integrated into the existing legacy network. This architecture faces the same issues as mentioned previously and common to other pre-encryption topologies. In addition, it experiences the additional cost burden of a legacy decryption element and the challenges of dynamically configuring and operating such a device. There may be additional costs faced in a specific deployment for switching and routing equipment that may be necessary to move transport streams "around" the legacy decryption device.

Changes are required to the method employed by the VOD system for creating dynamic PSI data to implement this architecture. The VOD system session manager must be aware of which conditional access method is appropriate for a session requested by a specific subscriber. This information must in turn be transferred to the video server that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The video server must be cognizant of the conditional access resources (ECMs) for each program stored on the server and these must be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, must indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

If one refers to the example movie scenario described in section 2.1, the same movie using 3.618GB of storage in the clear VOD state would require 3.618GBytes to store using re-encryption supporting two different CA systems.

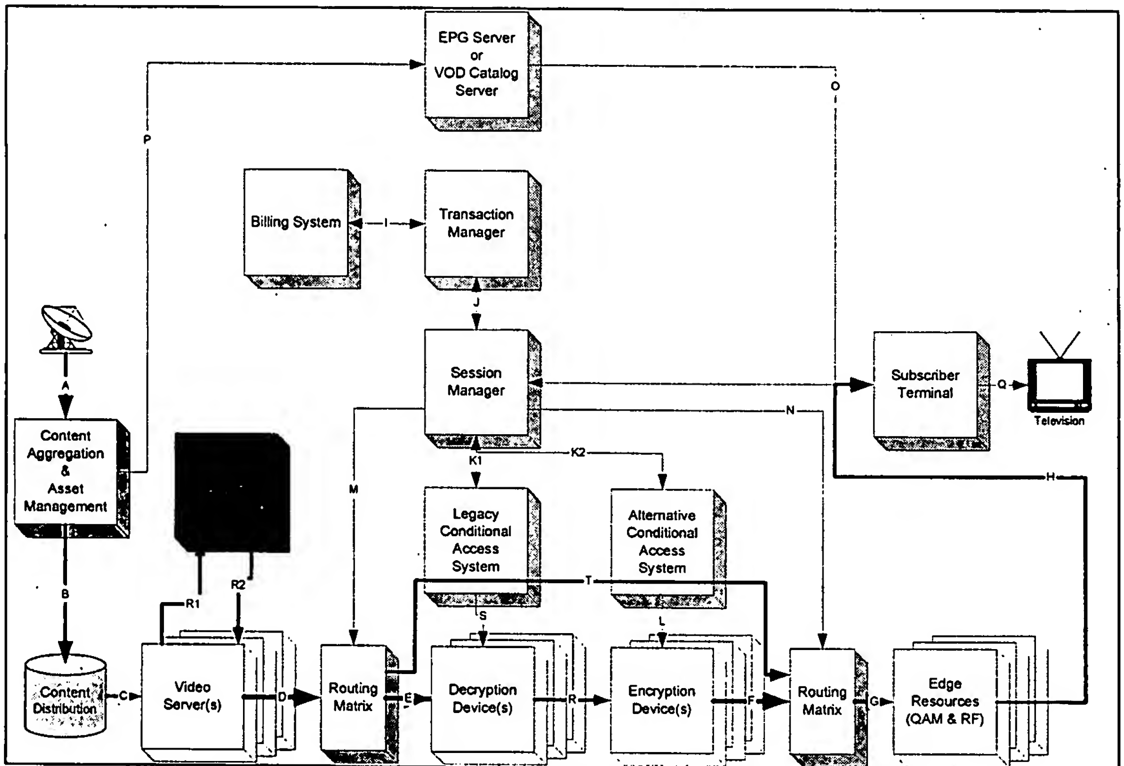

Figure 5 – *Re-Encrypted VOD Architecture*

2.2.5 Dynamic Composition Pre-Encryption

A final and more radical proposal for a pre-encrypted VOD architecture is dynamic composition pre-encryption. In this scheme, each program or movie is stored in three or more elements on the VOD video server. As can be seen in Figure 6, the content that is stored consists of either "critical" packets or non-critical packets. The "critical" packets constitute some 2% to 10% of the program and are encrypted. A separate copy of the critical content is maintained for each conditional access system supported by the MSO. The packets in both the "critical" packet file as well as the clear (unencrypted), non-critical packet file are indexed to maintain temporal correlation between the two files. These indices either may be monotonic packet counts from start of stream or calculated packet offsets from the last PCR.

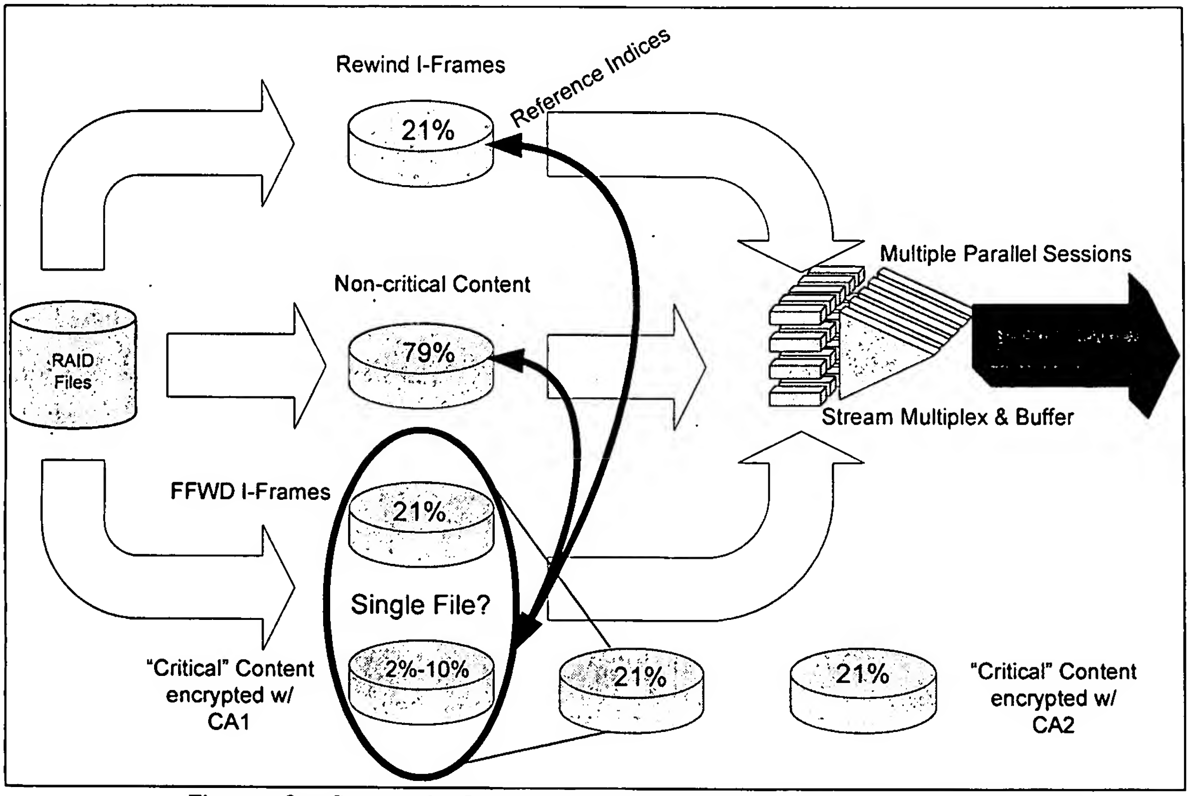

Figure 6 - *Dynamic Composition Pre-Encryption VOD Architecture*

When a subscriber session is initiated, the main file containing the clear content, less "critical" packets, is queued in the video server for playout. In addition, the file containing the "critical" packets, pre-encrypted in the CA format appropriate for the CPE of the subscriber requesting the session, is also queued for playout. When the program playback is started, the video server reconstructs a single program multiplex in its stream buffer feeding the outgoing transport the correct sequence of packets based upon the indices in the two component files. While the external composition and data flow appears identical to the clear VOD system depicted in Figure 1, the internal architecture of the video server changes significantly, as shown in Figure 6.

This method offers several distinct advantages that may not be readily apparent. The stream files containing "critical" packets may likely be the same one as the extracted subfile containing all I-frames for "trick" modes, as was described previously in the general discussion of VOD system architecture. If this opportunity is taken, then a storage economy can be realized over <u>all</u> pre-encrypted schemes including traditional (unencrypted) VOD, as deployed today. The traditional VOD video server has three files for each feature or movie: two containing just I-frames (one in reverse order) and one containing the complete original copy. Research on encoded streams conducted by Sony has shown that the I-frames typically represent between 12%-21% of the total content, typically around 17%. With the dynamic composition method, if the "critical" packet files are chosen to contain complete I-frames, the need for a separate file of critical data solely for encryption purposes is no longer necessary, saving 2% to 10% storage for this method. In addition, since this method removes the redundant I-frames from the clear stream file, an additional (nominal) 17% storage savings is also realized. This indicates a potential 27% (nominal, 31% maximum) video server disk storage savings for a single CA system model over the composite storage model VOD system described in section 2.2.2.

When compared to the segregated storage model described in section 2.2.1, one entire duplicate copy of a program is eliminated and the addition of <u>one</u> additional CA format adds **<u>NO</u>** storage or bandwidth overhead when compared to a traditional clear VOD server implementation. The reason for the "free" second CA format is that the 17% nominal storage saving realized by using the same I-frame file for both fast forward "trick" modes and "critical" content under the Passage™ scheme is consumed by replicating just the I-frame file and encrypting it with the alternative CA format.

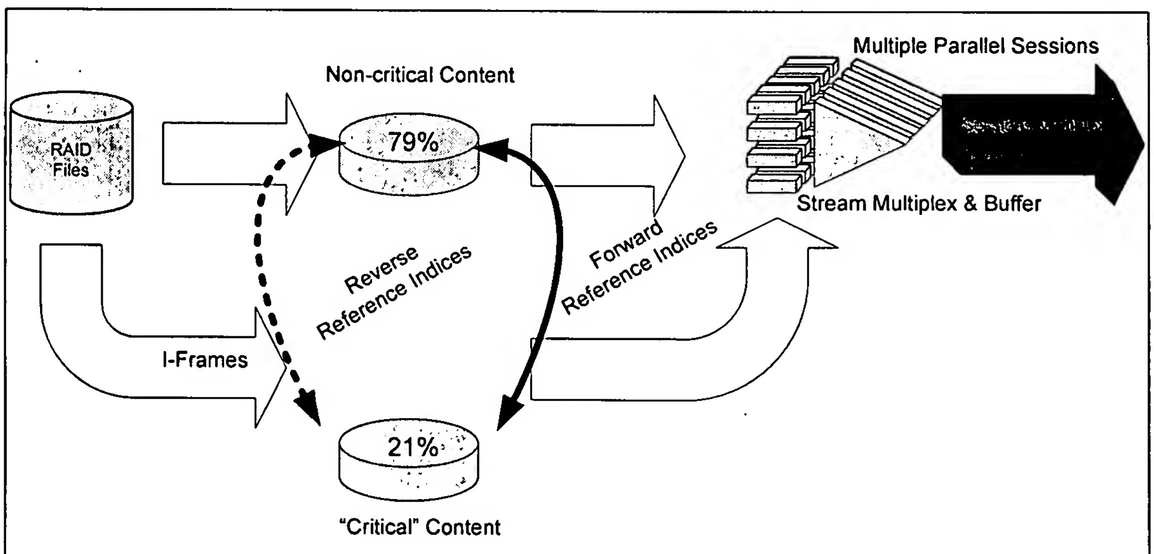

Figure 7 - *Dynamic Composition Pre-Encryption with Common Indices*

If one takes the concept one step further, the current convention in VOD systems to store the same I-frames of a movie in forward and reversed sequence to allow fast forward and rewind "trick" modes can be eliminated. It is replaced by one file of I-frames in normal forward sequence with two sets of indices, one for playing the I-frame file in forward order and one set for playing in reverse order. An illustration of this concept is shown in Figure 7. The appropriate sets of indices are chosen depending on whether forward or reverse high-speed motion is desired. The forward indices are also used to reconstruct the normal speed stream when matching the I-frame file to the non-critical content file to reconstruct the entire stream.. On a clear or re-encrypted VOD system, this will allow up to 21% storage savings. On a composite pre-encrypted storage system, up to 42% storage savings may be realized

Another advantage is that if the "trick" mode subfile and the "critical" data encrypted content file are the same, the content is selectively encrypted at a nominal 17% level, much higher than the commonly proposed Passage™ level of 2%, but carrying no inherent storage or system capacity costs, as do other schemes. For this system to work, some changes to the video server software design would be necessary, but these changes would be modifications to the existing processes and would not require substantial new development on the part of the server vendor.

A preprocessor would be required to perform selective encryption of content to be loaded onto the video server. This could simply be a reapplication of Passage™ processors currently under development by Passage™ licensees. A modified file protocol would be required to allow the video server to import and associate these files. Either the preprocessor or the video server could perform the indexing. An alternate instantiation would be to perform all Passage™ pre-processing within the video server itself. This could be accomplished by modifying the video server application of to add a pre-processor task as a separate executable, called by the server during the process to prepare content for pre-encryption. Obviously, these relationships need further clarification and codification in a series of system specifications.

Additionally, this method overcomes the classic pre-encryption issue of supporting trick modes, but retains the other common problems of encryption "shelf life" and the additional handling required to prepare the stream for use on the VOD system.

Changes are required to the method employed by the VOD system for creating dynamic PSI data to implement this architecture. The VOD system session manager must be aware of which conditional access method is appropriate for a session requested by a specific subscriber in order to select the appropriate "critical" data file for the session. This information must in turn be transferred to the video server that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The video server must be cognizant of the conditional access resources (ECMs) for each program stored on the server and these must be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, must indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

If one refers to the example movie scenario described in section 2.1, the same movie using 3.618GB of storage in the clear VOD state would require 3.159GBytes to store using dynamic composition pre-encryption supporting two different CA systems.

2.3 Session-Based Encryption VOD Distribution

Session based encryption is a widely discussed topic and there is much literature on the topic available. The basic premise is that a classic (clear) VOD server is modified to add an encryption device in series with the transport stream between the video server and the QAM modulator. In many systems, the encryption device may be integrated with the QAM modulator and other components. The Scientific-Atlanta MQAM and Harmonic NSG products are commercial examples of such devices.

The outgoing transport stream, containing multiple, independent VOD sessions and serving multiple subscribers, is encrypted at the point of distribution to the plant and in turn, subscribers. The control of the encryption and entitlements is based upon interaction between the session manager, which controls the session, video server and the conditional access system through defined interfaces. Most session based VOD architectures share the following common drawbacks:

- Coordination and/or distribution of entitlements and synchronization between session manager, conditional access system and stream encryption device.
- Security of the clear content from theft or piracy before loading on the video server and while stored in the system.
- Additional costs for adding both legacy and alternate stream encryption devices.
- Availability of legacy stream encryption devices with reasonable densities (session capacity).
- According to MSOs familiar with the subject, session based VOD streams are unsupported by present Motorola (former General Instrument) conditional access technologies.

One advantage of session-based encryption over the pre-encryption scheme is the additional security afforded by the application of unique encryption keys used for every session of the same program.

Another advantage is that in most cases, the video server does not need to generate special PSI that is aware of the conditional access method used for a specific session. The encryption device(s) downstream of the video server will append CA information specific to each session processed at the time/point of encryption. The VOD session manager still must manage which streams are processed by which CA method and in some cases, dynamically routing the streams to/through the encryption devices appropriate for a particular session.

As with other architectures, there are variations on the basic architecture of the session-based system and those variations are described below.

2.3.1 Segregated Session Based Encryption

Segregated session encryption is the extension of session-based encryption to multiple conditional access systems operating in conjunction with a single VOD system. It includes providing the appropriately encrypted stream for a specific subscriber session by routing the outgoing stream from the VOD video server to the subscriber on a transport stream & resultant RF carrier, carrying only a single common conditional access format. Sessions using other conditional access formats are similarly constrained (segregated) to other homogeneously encrypted transports/carriers. There is no sharing of resources between the CA systems and they operate independently.

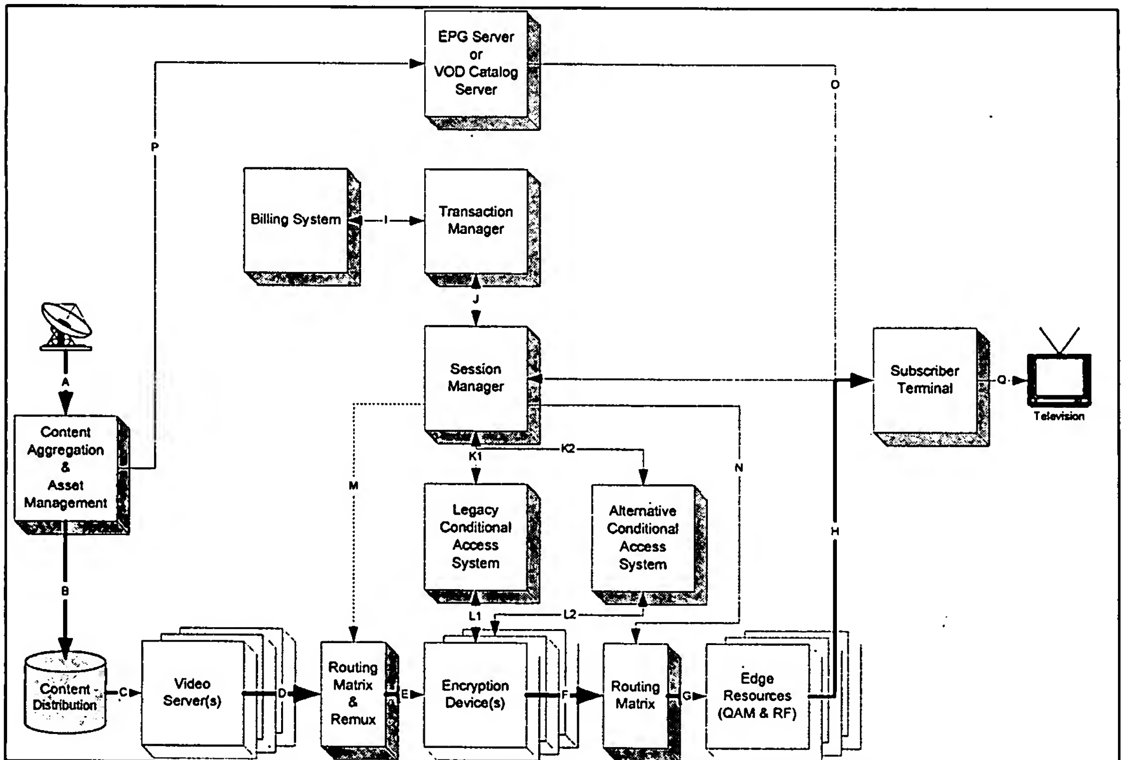

Figure 8 – *Segregated Session Based Encryption VOD Architecture*

During initiation of a new session, the VOD session manager determines which conditional access format is required by the requesting subscriber based upon information received either directly from the subscriber equipment or from another source, such as the billing system. The VOD session manager then determines the path to the appropriate encryption resource(s) having access to an RF node serving the subscriber's service area. This is done in a similar manner to the method used in large centralized clear VOD systems to find the appropriate video server(s) that can deliver a stream to the requesting subscriber.

Once a solution to the routing matrix is determined, the session manager coordinates the configuration of the routing elements and directs the CA system to apply encryption to the session through references to the assigned transport resources (PIDs).

This system presents a complex, real-time management requirement for determining usable resources available to apply to a new session and available spectrum transport slots. It requires equipment to perform stream routing (switch fabric) between the VOD video server(s) and the encryption elements, though these capabilities might be available integrated into other elements of the system. Additional spectrum is necessary to maintain segregation of the sessions on homogeneously encrypted transport streams and carriers.

A segregated session based encryption scheme requires, to some varying degree, duplication of encryption resources, since support of simultaneous sessions in differing conditional access formats is required. Careful traffic modeling is necessary to optimize the tradeoff between system capacity/resource availability and capital expenditure.

If one refers to the example movie scenario described in section 2.1, the same movie using 3.618GB of storage in the clear VOD state would require 3.618GBytes to store using segregated session based encryption supporting two different CA systems. The system could be optimized in a manner similar to that described in the section describing dynamic composition based pre-encryption. One I-frame file would be removed for rewind and a dual set of indices created for the remaining I-frame file to support both forward and reversed video sequences. In doing so, the total storage required for the example movie could be reduced to 3.159GBytes.

2.3.2 Composite Session Based Encryption

The composite session based encryption approach is similar to the segregated approach except that the transport streams/carriers provided to subscribers are heterogeneously encrypted. A single transport may contain any combination of two or more conditional access formats operating independently on a MPEG program basis, representing individual subscriber sessions.

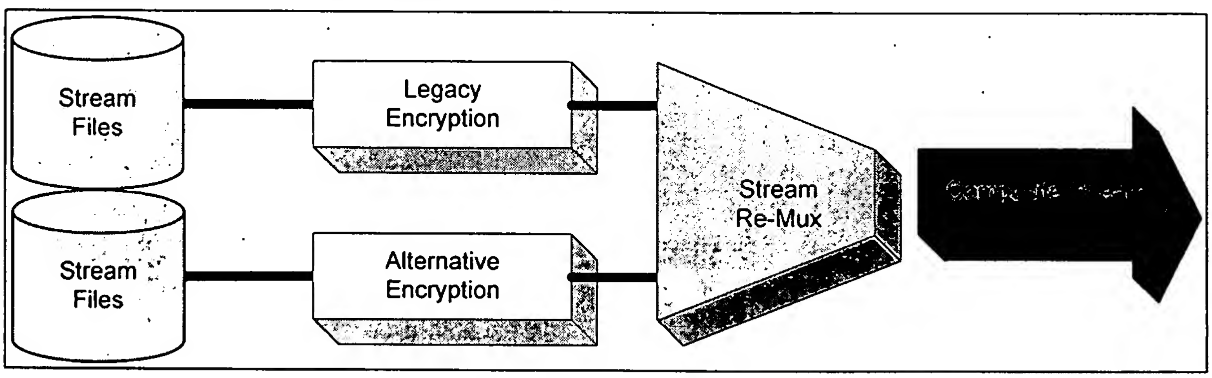

Figure 9 – *Composite Session Based Encryption Content Flow*

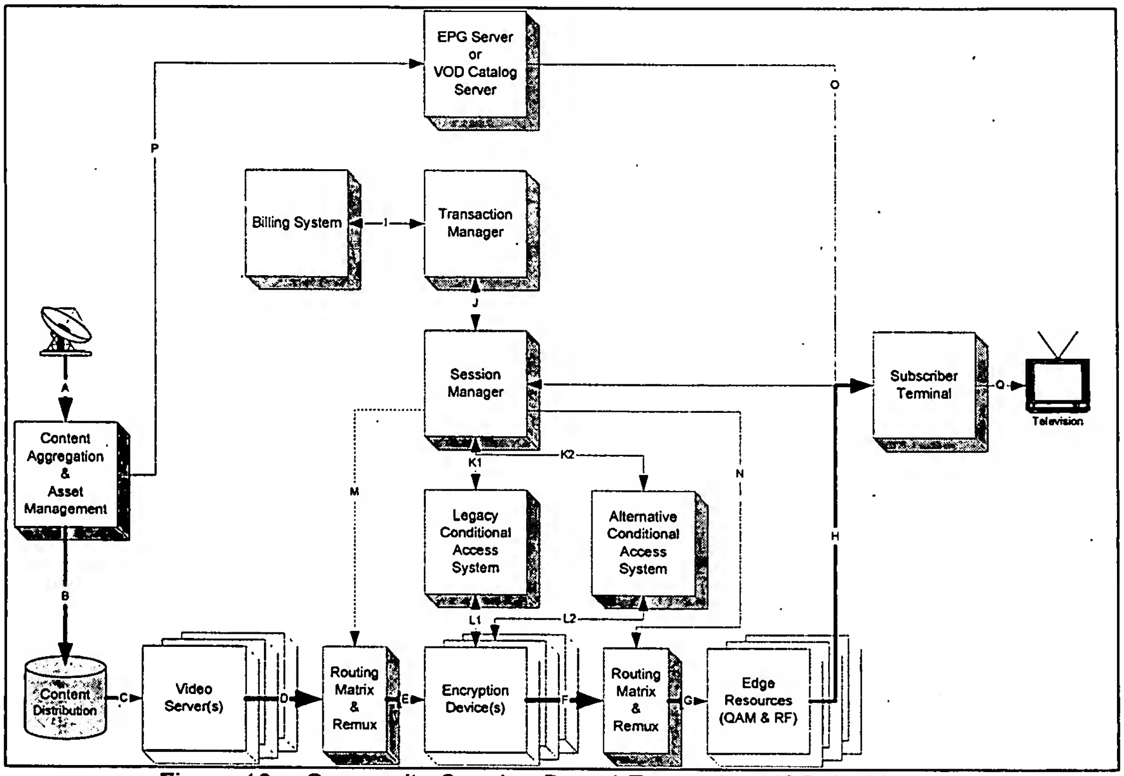

Figure 10 – *Composite Session Based Encryption VOD Architecture*

This scheme eliminates some of the complex real-time resource management processes required to determine available encryption resources, but instead trades it for the requirement that encryption resources appear in matched sets. Figure 9 and Figure 10 depict the design of a system employing composite, session based encryption.

The VOD session manager determines which CA format is appropriate for a given subscriber session and determines a VOD server that has access to the node representing the subscriber's service area. It then activates the appropriate CA resource in the encryption "set" attached to the node. An important point to note is that in no case is the Passage™ process of selective encryption employed, since there is never an opportunity to share any common content between subscriber sessions in a VOD paradigm. A technical consideration that may hinder integration is the configuration of systems with specific combinations of legacy encryption and/or remultiplexing equipment. This is especially true if the alternative encryption is embodied within the device performing the remultiplexing. The Harmonic NSG is a popular product for just this purpose. If the legacy system transmits data on unannounced PIDs or has critical latency concerns, this may be problematic if the device performing remultiplexing is not aware of these requirements.

If one refers to the example movie scenario described in section 2.1, the same movie using 3.618GB of storage in the clear VOD state would require 3.618GBytes to store using composite session based encryption supporting two different CA systems. The system could be optimized in a manner similar to that described in the section describing dynamic composition based pre-encryption. One I-frame file would be removed for rewind and a dual set of indices created for the remaining I-frame file to support both forward and reversed video sequences. In doing so, the total storage required for the example movie could be reduced to 3.159GBytes.

2.4 Batch-Based Encryption VOD Distribution

A revolutionary concept representing the best aspects of both session and pre-encrypted architectures is realized in batch-based encryption. As can be seen in Figure 11, the batch based VOD system has a topology different from the other systems presented in this document.

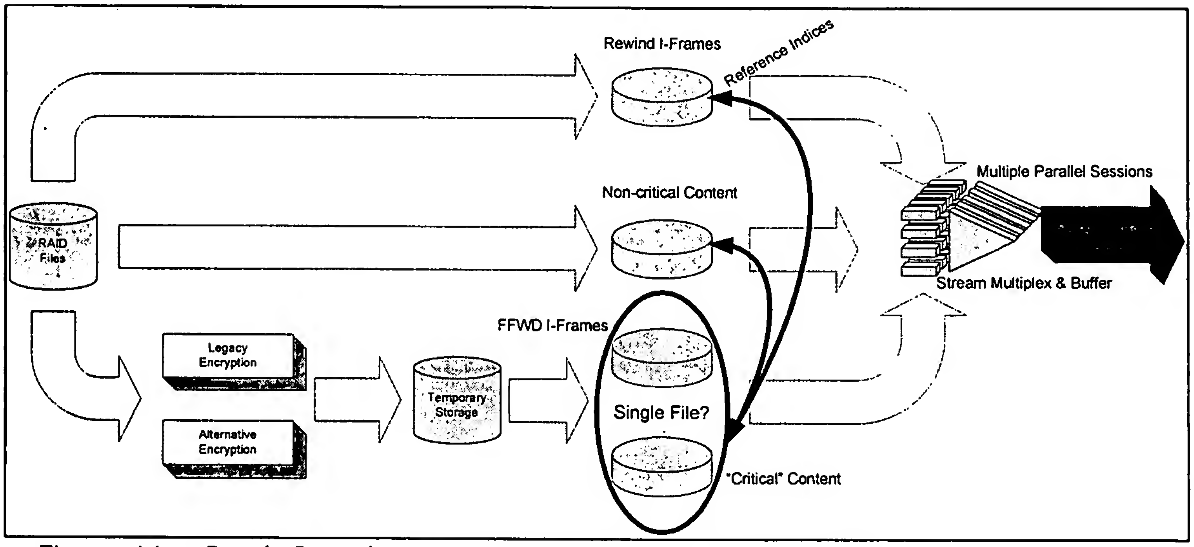

Figure 11 – *Batch-Based Encryption VOD Server Content Flow*

The content is stored entirely in the clear on the VOD video server, similar to the session-based system, but is contained in two files, representing "critical" packets and non-critical packets, just as in the case of the dynamic composition architecture. Likewise, the same opportunities for storage efficiency are available if the "critical" packet files are also used as the "trick" mode I-frame files. However, unlike the dynamic composition architecture, the "critical" packets are stored unencrypted. Additionally, this scheme departs from the dynamic composition architecture because there is no requirement to maintain an independent copy of the "critical" packet file for each conditional access system supported, providing further, substantial storage savings over the other architectures, typically on the order of 12% to 21% per conditional access system supported.

When the session manager initiates a new session at the request of a subscriber, the encryption technology appropriate for the subscriber's equipment is determined. The file for the selected feature containing the clear, non-critical content is queued in the video server for playout. In addition, a second file, containing the clear stream of "critical" packets is accessed; its contents are immediately streamed through a dedicated port on the video server at the maximum sustainable transport medium data rate (1Gbit/S for Gig-E, 200+Mbit/S for ASI, 38.8Mbit/S for DHEI) directly to the encryption resource identified by the session manager. This burst transferred file of I-frames, constituting only 12% to 21% of the video frames in the program is bulk encrypted at the highest rate that the encryption device and transport media can sustain. The encrypted I-frame content that emerges from the encryption device is captured to either a RAM or disk buffer resource within the VOD video server. For a typical 2-hour movie, with a nominal 17% I-frame content, this would require 450Mbytes of temporary storage per session.

When the program playback is started, the video server reconstructs in the stream buffer feeding the outgoing transport the correct sequence of packets based upon the indices in the clear, non-critical content component file and the smaller, batch-encrypted content that was captured back to the VOD video server, as described above.

This architecture, in addition to the storage efficiencies described both under the dynamic composition architecture description as well as in the previous paragraph, offers additional, significant advantages. The batch encryption of "critical" packet files allows for a significant reduction in the number of encryption devices required to provide encrypted delivery of VOD content. If one assumes support of two independent conditional access systems using this architecture, the I-frame and critical data residing in the same file and using a typical I-frame overhead (17%), then a single pair of encryption devices (incumbent & alternative) can support the same number of sessions as 60 pairs of encryption devices in a session based architecture (60:1).

An alternate instantiation could be to pre-encrypt sessions of I-frames and store them in the buffer for later consumption. In this manner, there would be no latency to delivering a new session due to the time overhead required to batch encrypt the file. The buffer of pre-encrypted I-frames could be replenished in the background to maintain a constant "inventory" of available sessions for delivery.

Changes are required to the method employed by the VOD system for creating dynamic PSI data to implement this architecture. The VOD system session manager must be aware of which conditional access method is appropriate for a session requested by a specific subscriber. This information must in turn be transferred to the video server that has been selected as the source for the session so that the appropriate PSI can be created for the session, including conditional access specific data. The video server must be cognizant of the conditional access resources (ECMs) for each program stored on the server and these must be dynamically allocated on unique PIDs along with PIDs for the corresponding audio and video data. The PSI generated for each specific session, in addition to indicating the assigned PIDs for A/V, must indicate the appropriate CASID, which is unique to each conditional access system provider and the PID assigned for the ECMs associated with the session.

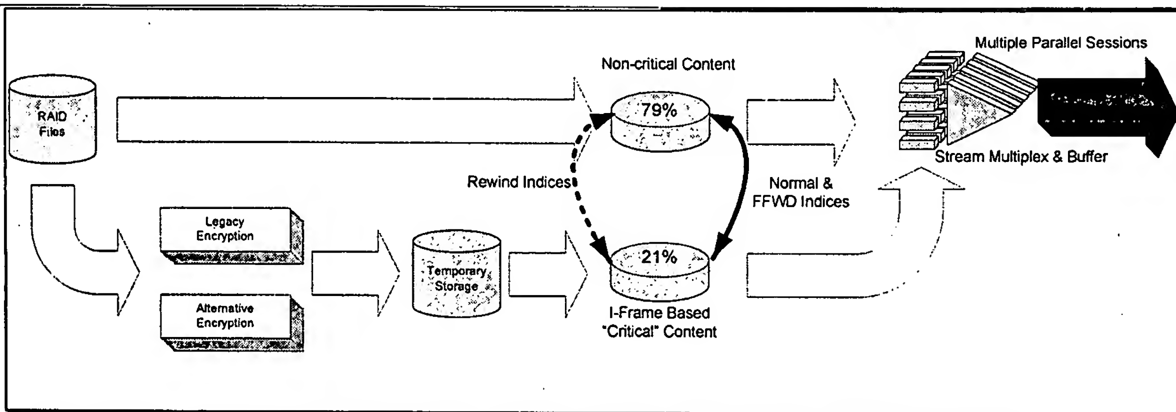

Figure 12 – *Optimized Batch-Based Encryption VOD Server Content Flow*

Another advantage over the pre-encryption scheme, shared with session-based encryption, is the additional security afforded by the application of unique encryption keys used for every session of the same program.

It is unclear at present whether this pseudo-pre-encryption format is compatible with Motorola systems.

If one refers to the example movie scenario described in section 2.1, the same movie using 3.618GB of storage in the clear VOD state would require 2.700GBytes to store using batch-based encryption supporting two different CA systems.

2.5 Pseudo-Simulcrypt VOD Distribution

The last architecture to be presented is the use of a simulcrypt-like process on a session-based encryption architecture. Simulcrypt provides for transport streams to be encrypted using a single, common encryption algorithm (DVB-CSA, DES, DES-ECB, 3-DES, AES, etc.) and the key to be delivered independently through entitlement control messages (ECMs) unique to each participating conditional access technology provider.

This scheme eliminates duplication of the incumbent/alternative encryption devices to process the transport streams from the VOD video server, which could allow realization of significant capital savings over other session-based schemes. The challenge is to determine whether the incumbent subscriber hardware supports activation of other, latent decryption algorithms. For example, a DCT-2000 may support other modes in addition to DC-2 (implemented on proprietary hardware, external to the demux/decoder IC) because it uses a commercially available decoder IC that does support alternative decryption formats directly on the IC. If the proprietary CA module can be bypassed or put into a passive, pass-through mode, the previously dormant decryption element in the decoder IC could be activated

If the platform is determined to support alternative encryption algorithms, then a business case must be made to gain cooperation of the incumbent. A technical challenge is the creation of ECMs that will cause the subscriber equipment to toggle between the current incumbent encryption algorithm that would be retained for compatibility with the broadcast transport system and the simulcrypt format intended to be used for VOD services.

3 Comparison Matrix

Table 1 – Summary Comparison of VOD Architectures

VOD Architecture	Relative Security	Additional Devices (Relative)	System Cost (Relative)	Operational Cost (Relative)	Scalability	Relative Storage Requirement	Incumbent Cooperation Required?
Clear VOD	Lowest	Lowest	Lowest	Lowest	Highest	Low:1	No
Segregated Storage	High:3	Low[1]:1	High	Highest	High:6	Highest:4	No
Composite Storage	Medium:1 (variable)	Low:2	Medium:2	High:7	High:6	High:3	No
Hybrid Composite Storage	Medium:1 (variable)	Low:2	Medium:2	High:7	Medium:4	High:3	No
Re-Encrypted Distribution	High:4	Medium:4	High	High:6	Low:3	Low:2	Yes[2]
Dynamic Composition	Medium:2	Low:2	Low:1	Low:5	High	Lowest	No
Segregated Session	Highest:5	High:5	High	Low:4	Low:1	Low:2	No
Composite Session	Highest:5	Highest[3]:6	High	Low:2	Lowest	Low:2	No
Batch Based Encryption	Medium:3	Medium:3	Medium:3	Low:1	High:5	Lowest	No
Simulcrypt	Medium	Medium	High	Low:3	Low:2	Low	Yes

[1] Assumes requirement for pre-encryption equipment
[2] Incumbent decryption devices, controllable by the VOD system is required
[3] One pair of incumbent and alternative encryption devices are required for each VOD transport stream

Table 2 - Summary Comparison of VOD Storage Requirements

Scenario: Video portion of a 2Hr movie@ 3Mbit/S CBR with 17% I-frame content, encrypted with 2 independent CA systems and as applicable, 2% Passage™ Critical Packet Density.

VOD Architecture	Storage	Optimized Storage[1]
Clear VOD	3.618GBytes	3.159GBytes
Segregated Storage	7.236GBytes	6.318GBytes
Composite Storage	3.690GBytes	3.222GBytes
Hybrid Composite Storage	3.690GBytes	3.222GBytes
Re-Encrypted Distribution	3.618GBytes	3.159GBytes
Dynamic Composition	3.159GBytes	N/A
Segregated Session	3.618GBytes	3.159GBytes
Composite Session	3.618GBytes	3.159GBytes
Batch Based Encryption	2.700GBytes	N/A
Simulcrypt	3.618GBytes	3.159GBytes

[1] Using method described in Section 2.2.5 for removal of redundant "trick" mode files.